

A Bayesian method for pulsar template generation

M. Imgrund^{1,2*}, D.J. Champion¹, M. Kramer^{1,3}, H. Lesch²

¹Max-Planck-Institute for Radio Astronomy, Auf dem Hügel 69, 53121 Bonn, Germany

²University Observatory Munich, Scheinerstraße 1, 81679 Munich, Germany

³Jodrell Bank Centre for Astrophysics, University of Manchester, Alan Turing Building, MP13 9PL, UK

Accepted x. Received x; in original form x.

ABSTRACT

Extracting Times of Arrival from pulsar radio signals depends on the knowledge of the pulsars pulse profile and how this template is generated. We examine pulsar template generation with Bayesian methods. We will contrast the classical generation mechanism of averaging intensity profiles with a new approach based on Bayesian inference. We introduce the Bayesian measurement model imposed and derive the algorithm to reconstruct a “statistical template” out of noisy data. The properties of these “statistical templates” are analysed both with simulated and real measurement data from PSR B1133+16. We explain how to put this new form of template to use in analysing secondary parameters of interest and give various examples:

We show how to reconstruct a detuned measurement’s phase shifts and demonstrate how to discriminate between different modes of radiation by implementing a nulling detection. Combining elements of the former, we implement a nonlinear filter for determining ToAs of pulsars. Applying this method to data from PSR J1713+0747 we derive ToAs self consistently, meaning all epochs were timed and we used the same epochs for template generation. The phase shift reconstruction is found to measure a shift in simulated data up to the estimated statistically possible accuracy out of noisy data.

Average templates as well as Bayesian templates are subject to uncertainties by fluctuations and noise. While the average template contains these as unavoidable artifacts, we find that the “statistical template” derived by Bayesian inference quantifies fluctuations and remaining uncertainty. This is why the algorithm suggested turns out to reconstruct templates of statistical significance from as few as ten to fifty single pulses.

A moving data window of fifty pulses, taking out one single pulse at the beginning and adding one at the end of the window unravels the characteristics of the methods to be compared. It shows that the change induced in the classical reconstruction is dominated by random fluctuations for the average template, while statistically significant changes drive the dynamics of the proposed method’s reconstruction.

The analysis of phase shifts with simulated data reveals that the proposed nonlinear algorithm is able to reconstruct correct phase information along with an acceptable estimation of the remaining uncertainty.

Key words: pulsars: general - methods: data analysis - methods: analytical - methods: statistical - methods: numerical - pulsars: individual (B1133+16, J1713+0747)

1 INTRODUCTION

Pulsars are well known for their high precision measurements, particularly in areas of fundamental physics such as general relativity (Cordes et al. 2004; Kramer et al. 2006). These measurements make use of the exceptional rotational stability of the pulsar as observed via its electromagnetic pulses, most often observed in the radio. Most of these analyses rely on determining the arrival times

of the pulses out of tens of thousands integrated single pulses.

The individual (or single-) pulses observed from a pulsar are extremely variable, not only in flux but also in the shape of the pulse. It is only when 10s of 1000s of pulses are averaged in an integrated pulse profile that it becomes stable (Helfand et al. 1975; Liu et al. 2012). Some pulsars exhibit temporal variations in this pulse profile; showing no emission for a period of time - nulling (see Sec. 3.5), or switching between one or more additional stable profiles (see Sec. 3.7), however for the majority of pulsars the pulse profile remains stable, in some cases over at least years).

* E-mail: mimgrund@mpifr-bonn.mpg.de

There exist ready to use toolsets like PSRCHIVE (Hotan et al. 2004) and TEMPO2 (Edwards et al. 2006; Hobbs et al. 2006) to analyse raw observation data in order to assign times of arrival to integrated observations and constrain the parameters of the physical model under consideration. The remaining differences between the model's prediction of an arrival time and the actual time measured, also known as timing residuals, can be as low as a few tens or hundreds nanoseconds on average while the periods of observational data spans several years. This amounts to tracking the rotational phase of the pulsar at a certain observational epoch to an accuracy of 10^{-4} and well below. In the light of a single pulse being a highly variable object that apparently cannot be tracked down to that accuracy (see e.g. fig. 11) it becomes clear that this very exact phase information is imprinted in and has to be recovered out of tens of thousands of very individual pulses.

The classical way of generating times of arrival (ToAs) is based on techniques comparing the shift of an epoch under consideration to a so called template. The template is often an integrated pulse profile of the entire dataset for that pulsar (at the appropriate observing frequency and bandwidth). Usually it is then fit with an analytic function to produce a noiseless template. These templates are occasionally updated and the ToAs reproduced.

Observations integrated over usually shorter times are then tested for a phase shift w.r.t. the template. The user of PSRCHIVE e.g. has the choice among five algorithms all comparing the data of an observation to the reference and yielding a relative shift and an error estimate on it for example by discrete cross-correlation and interpolation to yield sub-bin accuracy like Hotan et al. (2005) or examining the phase gradient in the Fourier space representation as in Taylor (1992) to name just a few. The phase shift is then converted into a time of arrival using the timestamp and folding period of the observation. In a next step, the actual pulsar timing is carried out. Pulsar timing is the process by which a rotational model of the pulsar is produced (the ephemeris) and the value of its rotation phase is mapped to the ToA at the observatory of every pulse. This is then fit to the ToAs from observations to produce a phase coherent solution that accounts for every rotation of the pulsar from the first observation to the most recent. It is the coherence of this solution that gives pulsar timing its extraordinary precision.

This constitutes the classical procedure of pulsar timing. Both the template and the observed profile to be converted to a ToA rely on integrated and thus averaged data before correlating. Single pulses cannot be compared directly for they are often not bright enough to be distinguished from noise introduced on the signal's way to the telescope and by the antenna-receiver-amplifier system. Furthermore their distinct and individual shape makes it difficult to correlate them to the average profile as simply matching the shape of the pulse against the average is not possible. Thus timing individual pulses by comparing them to a template lacks the necessary precision.

For these reasons and yielding a such precise timing the classical way of integrating the single pulses to yield a stable and confident profile justifies itself. However in actual observations, the reduced χ^2 value, a measure for the believed precision of the method compared to the actual residual error of the timing model, is much larger than unity (see e.g. Bailes (2010)) pointing to an underestimation of the error by the classical procedure and algorithms. This additional error is believed to have its root in additional noise sources like e.g. scattering by the interstellar medium. A comprehensive overview of the noise budgets can be found e.g. in Cordes & Shannon (2010). By assessing and fitting the possibly introduced effects on the ToAs, mitigating errors after determining ToAs has been inten-

sively investigated Coles et al. (2011); Cordes & Shannon (2010). As Bayesian methods already have been successfully used for correlating ToA-data from different pulsars, extending these to mitigate the timing errors by supposing a Bayesian timing model on the ToAs as described by Lee et al. (2014); Vigeland & Vallisneri (2013); Lentati et al. (2014) and others suggests itself.

The shape of single pulses is affected mainly by the short term noise contributions coming from current observational conditions one can aptly paraphrase as interstellar weather and the magnetosphere or radiation process intrinsic pulse jitter. Both influences affect mainly the single pulse statistics or influence the profile over short timescales, but can systematically influence timing and even dominate timing precision (Osłowski et al. 2011).

Increasing the gain of the observations with future systems like the Square Kilometre Array (SKA) aggravates this problem and demands for longer integration times if no solution is found to incorporate the single pulse fluctuations into the problem of determining ToAs rather than trying to average them out (Liu et al. 2012).

Finding a generic way of dealing with single pulse individuality can ameliorate the situation and fill the statistical gap between single pulses and an integrated profile without the need to take further modelling assumptions¹.

In this paper we will show a way to address the problem of generating pulsar ToAs with single pulse statistics and the benefits of having a more accurate statistical representation for single pulses' behaviour. By using more information than the classical method (which integrates this information) we hope to improve the precision of pulsar ToA generation and to allow ToAs from shorter observations to be unbiased by pulse jitter.

We introduce a model for statistics of the single pulses and substitute the classical template by a "statistical template" representing the single pulse statistics. We will then work out the steps of statistical template generation and phase shift detection replacing the classical counterparts.

Hereto we will derive a measurement model for squared and amplified receiver voltage in Sec. 2 and argue how to handle the signals from pulsars best, splitting them in a radiation process and a phase coupled template part. Using the Bayesian formalism, we will deduce the posterior distribution for both the amplitude and phase model introduced and thus show a way to infer a statistical template from input data. Though focusing on the important application of generating ToAs, we will also demonstrate the benefit of using statistical templates on other applications such as determining nulling and analysing a moding pulsar. Sec. 3 will evaluate the proposed algorithm and methods using both simulated and real pulsar data also giving insight on how a statistical template can enhance the view on the pulsars profile from a statistical perspective. Comparing the generation of ToAs to the classical results we will highlight improvements both in timing residuals as well as the χ^2 value of the fit. We summarize our findings in Sec. 4.

2 MODELLING PULSAR MEASUREMENTS

2.1 Terminology of the used statistical model for single pulses

In order to describe single pulse statistics we use the notion of a fluctuating envelope besides to the classical set of terms of single pulses and integrated profiles.

¹ Incorporating assumptions on e.g. scattering of the ISM or pulse jitter behaviour however may further improve the results presented




Figure 1. Scheme for single pulse statistics as described in Sec. 2.1

Instead of modeling an average profile, we assume a single pulse to be the result of multiplying a stationary Gaussian random noise process, which amounts to the physical process that generates the radiation in the pulsar magnetosphere, with a possibly unstable envelope (see fig.1). This envelope models the periodic influence of the magnetosphere. The uncertainty in the envelope may reduce with observation time when the envelope is not fluctuating. However if the magnetosphere conditions for some region mapping to the pulse phase are fluctuating, the envelope will stay fluctuating in that region.

Only in the limit of no uncertainty the envelope would equal the (scaled) average of the pulses and amount to the same thing as a classical integrated profile. As averaging single pulses generates a pulse profile, we emphasize that there are many possible fluctuating envelopes which are integrated to the same profile.

The empiric or analytic template that is used for generating times of arrival will consequently be substituted by a Bayesian statistical template describing the fluctuating envelope.

2.2 Amplitude model

Every Bayesian analysis is based on formulating a signal and noise model to describe the statistical imprint of a certain signal in the data. Thus, before inserting the outlined statistical model of the signal, we have to understand the influence of an arbitrary signal. Therefore we will introduce the basic measurement model, insert our signal model and then derive the formulae for signal reconstruction using Bayes' theorem.

Typical pulsar measurements at radio frequencies give the intensity of received radiation by suitable sampling of \mathbf{d} , the quadrature of the received voltage. The problem is to infer a unknown signal \mathbf{s} (which in our case amounts to the noise-free series of single pulses as emitted from the pulsar) from the noisy data \mathbf{d} . The noise - leaving aside radio interference - is dominated by thermal fluctuations of both receiver and antenna and is accurately described by white noise of a usually known variance σ_n^2 entering before quadrature of the signal. Thus, the baseline corrected data \mathbf{d} measured, given a signal \mathbf{s} and a random noise \mathbf{n} amounts to

$$\mathbf{d} = (\mathbf{n} + \mathbf{s})^2 \quad (1)$$

$$\mathbf{d} - \mathbf{n}^2 = \mathbf{s}^2 + 2\mathbf{s}\mathbf{n} \quad (2)$$

where

$$\mathcal{P}(\mathbf{d}|\mathbf{s}, \mathbf{n}) = \delta(\mathbf{d} - (\mathbf{n} + \mathbf{s})^2) \quad (3)$$

$$\mathcal{P}(\mathbf{n}) = \mathcal{G}_{\mathbf{n}}(0, \mathbf{N}) \quad (4)$$

where $\mathcal{P}(x|y)$ is the short notation for the probability density of $X = x$ given $Y = y$, where X, Y would be the corresponding random variables. We use boldface lower case letters to emphasize that data, signal and noise consist of many single values $\mathbf{d} = \{d_1, d_2, \dots, d_t, \dots\}$ for which the equations stated are valid component-wise. Thus operations like taking a root of a vector or multiplying two vectors are carried out only on the components of the vector. For the scalar product of vectors \mathbf{a} and \mathbf{b} we use the notations $\mathbf{a}^\dagger \mathbf{b}$ and $\sqrt{\mathbf{a}^\dagger \mathbf{a}} = \|\mathbf{a}\|$. We abbreviate $\det(\cdot)$ as $|\cdot|$ and define $\delta(\mathbf{a}) := \prod_k \delta(a_k)$. Putting a hat over a vector $\hat{\mathbf{v}} = \text{diagmat}(\mathbf{v})$ denotes constructing a diagonal matrix with a diagonal with the elements of \mathbf{v} . When dealing with sets, $\mathbf{A} \setminus \mathbf{g}$ denotes the set \mathbf{A} without the elements \mathbf{g} .

$\mathcal{G}_{\mathbf{x}}(\mathbf{m}, \mathbf{V})$ denotes that \mathbf{x} is distributed as the multivariate Gaussian distribution with mean \mathbf{m} and covariance matrix \mathbf{V} . It is defined as

$$\mathcal{G}_{\mathbf{x}}(\mathbf{m}, \mathbf{V}) := \frac{1}{\sqrt{2\pi|\mathbf{V}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \mathbf{m})^\dagger \mathbf{V}^{-1}(\mathbf{x} - \mathbf{m})\right]$$

Integrating over a vector \mathbf{a} is denoted by $\mathcal{P}(\mathbf{b}) = \iint \mathbf{D}\mathbf{a} \mathcal{P}(\mathbf{a}, \mathbf{b}) = \iint da_1 \dots da_k \mathcal{P}(\mathbf{a}, \mathbf{b})$.

The classical way of determining the template intensity amounts to averaging eq.(2) over several pulsar periods, noticing that the term mixing signal and noise vanishes leaving us with the desired quantity of $\langle s^2 \rangle$ when subtracting σ_n from the averaged data. To use Bayes' theorem, we will instead calculate the exact probability distribution to measure \mathbf{d} at a certain time, given a certain signal value \mathbf{s} at that time:

$$\begin{aligned} \mathcal{P}(\mathbf{d}|\mathbf{s}) &= \iint \mathbf{D}\mathbf{n} \mathcal{P}(\mathbf{n}, \mathbf{d}|\mathbf{s}) = \iint \mathbf{D}\mathbf{n} \mathcal{P}(\mathbf{d}|\mathbf{n}, \dots) \mathcal{P}(\mathbf{n}) \\ &= \iint \mathbf{D}\mathbf{n} \delta(\mathbf{d} - (\mathbf{n} + \mathbf{s})^2) \frac{1}{\sqrt{|\mathbf{N}|}} \exp\left[-\frac{1}{2} \frac{\|\mathbf{n}\|^2}{\sigma_n^2}\right] \end{aligned} \quad (5)$$

We substitute $\tilde{\mathbf{n}} = (\mathbf{n} + \mathbf{s})^2$, $\mathbf{n} = \pm\sqrt{\tilde{\mathbf{n}}} - \mathbf{s}$, $d\mathbf{n} = \frac{d\tilde{\mathbf{n}}}{2\sqrt{\tilde{\mathbf{n}}}}$ and notice that $\sqrt{|\mathbf{N}|} = \sigma_n^{N_{\text{tot}}}$, where N_{tot} is the total number of samples in all bins.

$$= \int \int_0^\infty \mathbf{D} \frac{\tilde{\mathbf{n}}}{\sqrt{\tilde{\mathbf{n}}}} \delta(\mathbf{d} - \tilde{\mathbf{n}}) \frac{1}{\sqrt{2\pi\sigma_n^{N_{\text{tot}}}}} \exp\left[-\frac{1}{2} \frac{\|\sqrt{\tilde{\mathbf{n}}} \pm \mathbf{s}\|^2}{\sigma_n^2}\right] \quad (7)$$

$$= \frac{1}{\sqrt{2\pi\mathbf{d}}\sigma_n^{N_{\text{tot}}}} \exp\left[-\frac{1}{2} \frac{\|\sqrt{\mathbf{d}} \pm \mathbf{s}\|^2}{\sigma_n^2}\right] \quad (8)$$

where we now have to assign a signal model for \mathbf{s} . We assume

$$\mathbf{s} = \exp[\mathbf{f}] \cdot \mathbf{g}\sigma_n \quad (9)$$

Where we have decided for the signal to be measured in units of noise temperature of the receiver-antenna system. The \pm -operator means that the whole density would be a sum of the densities with the respective sign. This operator vanishes later on, when integrating over \mathbf{g} and using $\mathcal{P}(\mathbf{g}) = \mathcal{P}(-\mathbf{g})$. \mathbf{f} and \mathbf{g} are Gaussian random variables with also to be determined covariances \mathbf{F} and \mathbf{G} which describe the statistical properties of the envelope respectively the radiation process. The exponential form for the envelope was chosen since on one hand it is able to grasp the occurring fluctuations to high radiation densities of single pulses and on the other we may

easily assume a scale-free prior for the amplitude, as we will see below.

We assume the radiation process to have stationary statistics given by

$$\mathbf{G} = \mathbf{G}(t - t')$$

and the part that is generating the observed envelope to be periodic in time. Even though the covariance matrix of $\exp[f]$ alone has non-diagonal components, the product's covariance matrix $\langle g(t)\exp[f](t)g'(t)\exp[f]'(t) \rangle$ has only diagonal components if we marginalize over g , because of the stationarity of g as assumed in (2.2). Thus we describe the covariance matrix solely by its Fourier-coefficients at $\omega_k = k \cdot \frac{2\pi}{T}$. Since the expectation of the signal's covariance matrix for a certain g integrated over a specific pulse certainly does not vanish, the non-diagonal components of the signal's matrix carry information about g which we will discard in this paper. T denotes the real pulsar period, which we will discriminate from the assumed period τ later on.

2.2.1 Joint Probability and Information Hamiltonian

For a moment, let us assume the covariance matrices \mathbf{F} and \mathbf{G} are known. Then we may write for the joint probability

$$P(\underbrace{\mathbf{f}, \mathbf{g}, \mathbf{s}, \tau}_{=: A \cup \{s\}} | \underbrace{\mathbf{F}, \mathbf{G}}_{=: B}) = \delta(\exp[\mathbf{f}]\mathbf{g}\sigma_n - \mathbf{s})\mathcal{G}_{\mathbf{f}}(0, \mathbf{F})\mathcal{G}_{\mathbf{g}}(0, \mathbf{G})\mathcal{P}(\tau)$$

and we define the parameter sets A and B where $\mathcal{P}(\tau)$ is the probability density over the assumed period τ . In order to define the Information Hamiltonian $H_B[A, \mathbf{d}] := -\log P(A, \mathbf{d}|B)$ we notice that $\mathcal{P}(A, \mathbf{d}|B) = \iint \mathcal{D}\mathbf{s} \mathcal{P}(\mathbf{d}|\mathbf{s})P(A \cup \{s\}|B)$ and use the measurement model given by (8):

$$\begin{aligned} -\log P(A, \mathbf{d}|B) &= H_B[A] = \\ &= \frac{1}{2} \left[(\sqrt{\mathbf{d}} - \exp[\mathbf{f}] \cdot \mathbf{g}\sigma_n)^\dagger \frac{1}{\sigma_n^2} (\sqrt{\mathbf{d}} - \exp[\mathbf{f}] \cdot \mathbf{g}\sigma_n) + \right. \\ &\quad \left. \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f} + \mathbf{g}^\dagger \mathbf{G}^{-1} \mathbf{g} \right] - \\ &\quad - \log \mathcal{P}(\tau) + \frac{1}{2} \log(|2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |2\pi\mathbf{G}| |\mathbf{d}|) \end{aligned} \quad (10)$$

Since we are interested in the pulsar's statistical template given by the envelope characteristics we marginalize over the radiation process g by integrating it out. This calculation is outlined in Sec. A1. We are left with

$$\begin{aligned} H_B[A] &= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} [1 + \widehat{\exp[\mathbf{f}]\mathbf{G}\exp[\mathbf{f}]}]^{-1} \sqrt{\mathbf{d}} + \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f} + \right. \\ &\quad \left. + \log(|2\pi\mathbf{D}_{\mathbf{f}}^{-1}| |2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |2\pi\mathbf{G}| |\mathbf{d}|) \right] - \log \mathcal{P}(\tau) \end{aligned} \quad (11)$$

where the operator $[1 + \widehat{\exp[\mathbf{f}]\mathbf{G}\exp[\mathbf{f}]}]^{-1}$ together with the noise variance is a variation of the operator leading to the well-known case of the Wiener filter (Wiener 1949). This similarity arises due to the integration over the Gaussian random field \mathbf{g} added upon the signal field \mathbf{s} . However, the normalisation w.r.t. the data differs and we further distinguish between the periodic part described by (\mathbf{f}, \mathbf{F}) and the radiation part now showing up only as \mathbf{G} . The functioning of the operator can be understood by considering its interplay with given data. Maximizing probability equals minimizing the Hamiltonian. Thus high data values will be compensated by large mean values of \mathbf{f} . Conversely, for too high values of \mathbf{f} the exact

value of the data becomes irrelevant, leading to a sub-optimal solution independent of the data. The optimal solution resembles the actual probability distribution the data would be drawn from. As a function of \mathbf{d} , the joint probability is a special case of a gamma distribution $\propto \exp[-|\frac{\mathbf{d}}{\mathbf{m}}|] |\mathbf{d}|^{-.5}$ where \mathbf{m} is the mean value of the distribution. In this light the operator resembles the classical formula $\langle \mathbf{d} \rangle = \mathbf{m} = \sigma_n^2 + \sigma_s^2$ since $\mathbf{m} \hat{=} \sigma_n^2 [1 + \frac{\mathbf{s}^\dagger \mathbf{s}}{\sigma_n^2}]$ for our case. \mathbf{G} becomes diagonal in Fourier space since \mathbf{g} is assumed to be stationary. As $\exp[\mathbf{f}]$ is a periodic signal, its Fourier space is limited to discrete frequencies that are a multiple of the assumed periodicity. We may absorb the values of G_ω , where $\omega_k = k \cdot \frac{2\pi}{\tau}$ into \mathbf{F} . In the following we will discard information about \mathbf{G} from single pulses. Since our main interest lies on the statistics of the envelope, we may take this loss.

Under these assumptions, and neglecting non-diagonal terms arising when Fourier transforming f (see Sec. B and 2.2.3), the $G_\omega =: \sigma_{g_k}^2$ may be absorbed into an effective mean \mathbf{f} as is evident in eq. (11):

$$\exp[2f_k] \cdot \sigma_{g_k}^2 = \exp[2(f_k + \bar{f}_k)]$$

where $\bar{f}_k = \log \sigma_{g_k}$. We redefine \mathbf{f} in (11) in that way and yield

$$\begin{aligned} H_B[A \setminus \mathbf{g}] &= \frac{1}{2} \left[\left\| \frac{\mathbf{d}}{\sigma_n^2} \frac{1}{1 + \exp[2\mathbf{f}]} \right\|^2 + (\mathbf{f} - \bar{\mathbf{f}})^\dagger \mathbf{F}^{-1} (\mathbf{f} - \bar{\mathbf{f}}) + \right. \\ &\quad \left. + \log(|2\pi\mathbf{D}_{\mathbf{f}}^{-1}| |2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |\mathbf{d}|) \right] - \log \mathcal{P}(\tau) \end{aligned} \quad (12)$$

where \mathbf{G} is now the unity matrix and we assume \mathbf{F}^{-1} to be diagonal with components $\frac{1}{\sigma_{f_k}^2}$. We arrive at our model for a single pulse's amplitudes in Fourier space, assuming all timing parameters to be given:

$$\begin{aligned} H_{\bar{\mathbf{f}}, \sigma_f^2}[\mathbf{f}, \mathbf{d}] &= \frac{1}{2} \left[\left\| \frac{\mathbf{d}}{\sigma_n^2} \frac{1}{1 + \exp[2\mathbf{f}]} \right\|^2 + (\mathbf{f} - \bar{\mathbf{f}})^\dagger \frac{1}{\sigma_f^2} (\mathbf{f} - \bar{\mathbf{f}}) + \right. \\ &\quad \left. + \log C(\mathbf{F}, \sigma_n, \mathbf{d}) \right] \end{aligned} \quad (13)$$

This is the negative logarithm of the likelihood function. Bayes' law gives us the posterior

$$P(\mathbf{f}|\mathbf{d}) = \frac{\exp[-H_{\bar{\mathbf{f}}, \sigma_f^2}[\mathbf{f}, \mathbf{d}]]}{\iint \mathcal{D}\mathbf{f} \exp[-H_{\bar{\mathbf{f}}, \sigma_f^2}[\mathbf{f}, \mathbf{d}]]} \quad (14)$$

when we assume a flat prior on $\bar{\mathbf{f}}$ which effects in a scale-free prior on the assumed amplitude, since $\mathbf{s} \propto \exp[\mathbf{f}]$ is exponentiated.

2.2.2 The receiver equation's pdf

Assuming G to be a diagonal of ones (meaning vanishing correlations) for this argument, the pdf of the receiver equation of each bin decouples and for a specific bin and different pulses K takes the form of

$$\mathcal{P}(f|\vec{d}) \propto \prod_K \exp\left[-\frac{1}{2} \frac{d_K}{1 + f^2}\right] \sqrt{\frac{1}{(1 + f^2)2\pi}} / \sqrt{d_K}$$

which simply states that the variance of $\sqrt{d_K}$ can be explained as sum of the variance of the noise (the one in the denominator of the exponentiated term) plus the variance of g (equals unity per definition) times the signal strength of the envelope f . The square root term of d_K after the exponential is a normalization factor w.r.t. f . We deduce that for a constant signal envelope f and vanishing




Figure 2. Fourier transform of a whole epoch of pulsar B0329+54. The x axis has been scaled to the pulsar’s signal harmonics. Inset: detail of one spike.




Figure 3. Comparison of absolute and relative error (as uncertainty over actual signal strength) for a single measurement of a signal f

correlations of g

$$\begin{aligned} \mathcal{P}(f|\vec{d}) &\propto \exp\left[-\frac{1}{2} \frac{\sum_K d_K}{1+f^2}\right] \sqrt{\frac{1}{(1+f^2)}} \\ &= \exp\left[-\frac{1}{2} \frac{\langle d \rangle}{\frac{(1+f^2)}{N_K}}\right] \sqrt{\frac{1}{(1+f^2)^{N_K}}} \end{aligned} \quad (15)$$

where N_K is the number of pulses involved. This probability density function attains its maximum for $\langle d \rangle \geq 1$ at $f = \sqrt{\langle d \rangle - 1}$ justifying the average data subtracting the noise background as the maximum likely guess. Carrying out a saddle point analysis, we conclude that

$$\sigma_f = \sqrt{\frac{(1+f^2)^2}{(2N_K f^2)}} = \begin{cases} \sqrt{\frac{1}{2N_K}} \frac{1}{f} & \text{for } f \ll 1 \\ \sqrt{\frac{1}{2N_K}} f & \text{for } f \gg 1 \end{cases}$$

For a signal that is constant and not the modulation of a stationary noise process, we would have expected $\sigma_c \propto \sqrt{c/N_K}$ as $c \gg 1$ such that the relative error decreases with signal strength. In our case the relative error decreases too, but only to reach 1 as a limit (see Fig. 3). Against our intuition we expect the shape of the profile

to be most uncertain at the highest signal values while the areas of lowest uncertainty can be found where $f \approx 1$.

We emphasize that this relative uncertainty will not be integrated out by measuring more pulses, as the uncertainty in both high and low signal areas equally fall with the square root of N_K leaving the ratio of the remaining error untouched. This fact explains the often expressed observation that stronger pulses do exhibit larger timing errors (Shannon et al. 2014; Osłowski et al. 2011). The reason was found implicitly by numerous authors (e.g. Kulkarni (1989); Rickett (1975)) for the case of self-correlation of the signal. We emphasize that even without self-correlation, we expect the higher pulses to have a higher intrinsic uncertainty.

The non-diagonal terms introduced by G mark the departure from simply averaging the data. Let us review the influence of a non-white stationary process on the data’s spectrum.

2.2.3 The influence of g on the signal’s spectrum

When investigating a Fourier transform of a whole observation’s time series containing several hundred pulses (see fig. 2) two prominent features dictate the spectrum: Spikes of the pulsar’s signal dominate over receiver background noise. If the signal was strictly periodic, we would expect it to be described solely in terms of the Fourier coefficients of harmonics of the base frequency $1/P$. As the pulsar’s signal contains periodicity, we expect spikes to show up in this case too². However, we are facing a stochastic process g that is modulated via a periodic envelope f . As is commonly known, multiplication in time domain translates to convolution in Fourier domain. As $g(\omega)$ is expected to be a red noise process, we expect it to cast the sharp pulsar signal to finite width in Fourier space. Indeed - as the inlay of the figure shows - the signal’s power is leaked to the neighbouring bins. The width of the spectral line however is rather narrow, pointing to a fast decline of the spectrum

² these spikes are not an artifact of the folding period for the archive as this folding is discarded in the analysis. Of course the folding period of observations is set as close to the true period P as possible to be able to integrate the data easily by adding up the receiver bins’ values (in soft- or hardware)

of g in Fourier space. Consequently, the spectrum of g in time domain is spread out widely, pointing to non-vanishing correlations over a long time. These correlations diminish the independence of both adjacent lying bins and the bins of neighbouring pulses.

As the multiplication $f(t) \cdot g(t)$ amounts to a convolution in Fourier space, evaluating the posterior distribution in Fourier space exactly was found to be unfeasible. Instead, we approximated the convolution by taking only the diagonal terms, losing mathematical rigour and precision but keeping the problem solvable in reasonable computational time. The interested reader may refer to Sec. B of the appendix for a more detailed discussion why solving this in time domain would be favorable, but currently is not feasible.

The clear departure from the average comes from the off-diagonal elements of $\sqrt{d}M\sqrt{d}^\dagger$ where M is the operator derived. As is evident from the structure of M , it is expected to be off-diagonal only until the correlations introduced by $G(t - t')$ decay. Thus, writing this equation as a matrix equation, symmetry tells us that we can possibly store only a limited combination of $\sqrt{d}(t - \Delta t)\sqrt{d}^*(t + \Delta t)$ values and average them as they are multiplied by the same number when performing the matrix multiplication. The matrix M , subject to the exact envelope shape and correlations of the radiation process can be calculated afterwards and former data may be re-processed with the new M in that way without losing the statistical information from the single pulse level, even though we may integrate parts of the single pulse data to yield smaller data sets. This could possibly be exploited in the future to profit from single pulse statistics without having to store all single pulses.

2.2.4 Maximum-a-posteriori-Filter

We are interested in finding the \bar{f}, σ_f values most compatible with the data and thus set out to maximize the a-posteriori distribution³ given the data. For a prior specified by $\sigma_{f,p}$ and \bar{f}_p we derive the integral over f of (14) w.r.t. \bar{f} and expect the derivative to vanish. This procedure yields an implicit formula for maximizing the posterior distribution of \bar{f} .

$$\begin{aligned} 0 &\stackrel{!}{=} \iint \text{Df} \sum_{K=0}^{K_{\max}} \sum_{k=0}^{k_{\max}} \frac{(f_{K,k} - \bar{f}_k)}{\sigma_k^2} \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \\ 0 &= \sum_{K=0}^{K_{\max}} \sum_{k=0}^{k_{\max}} \int \text{d}f_{K,k} \frac{(f_{K,k} - \bar{f}_k)}{\sigma_k^2} \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \end{aligned}$$

which may be solved for every k independently as

$$\begin{aligned} 0 &= \sum_{K=0}^{K_{\max}} \langle f_{K,k} - \bar{f}_k \rangle \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \\ \bar{f}_k &= \langle f_{K,k} \rangle \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \end{aligned} \quad (16)$$

where we organize our data in pulse numbers indexed K and have summed over single Fourier coefficients d_k of these data of pulse K .

The new variance assumed of σ_f may be calculated equally by evaluating the expectation value of $(f_{k,K} - \bar{f}_k)^2$, since deriving w.r.t.

³ Maximum A Posteriori or MAP refers to taking the most likely value of the posterior as best guess

σ_k yields:

$$\begin{aligned} 0 &\stackrel{!}{=} \int \text{d}f \sum_{K=0}^{K_{\max}} \left(\frac{(f_{K,k} - \bar{f}_k)^2}{\sigma_k^3} - \frac{1}{\sigma_k} \right) \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \\ 0 &= \sum_{K=0}^{K_{\max}} \sum_{k=0}^{k_{\max}} \int \text{d}f_{K,k} \left(\frac{(f_{K,k} - \bar{f}_k)^2}{\sigma_k^3} - \frac{1}{\sigma_k} \right) \exp[\dots] \end{aligned}$$

which may be solved for every k independently as

$$\begin{aligned} 0 &= \sum_{K=0}^{K_{\max}} \langle ((f_{K,k} - \bar{f}_k)^2 - \sigma_k^2) \rangle \exp[-H_{\bar{f}_k, p, \sigma_{f_p}^2} [f_{K,k}, d_{K,k}]] \\ \sigma_k &= \sqrt{\langle (f_{K,k} - \bar{f}_k)^2 \rangle \exp[\dots]} \end{aligned} \quad (17)$$

Doing this once is assuming a prior of the form of the initial \bar{f}, σ_f . A very general choice may be suitable. Then taking the expectancy values of (16) as prior input to (14) and iterating is known as the expectation-maximization (EM) algorithm proposed by Dempster et al. (1977). We may in this way in principle get rid of assumptions going into the first prior by iterating since this procedure converges on a prior fully compliant with the data, assuming gaussian distribution of f . However the EM algorithm may only find local maximum which equals a global maximum if the pulsar is really radiating with a log-normal probability distribution. The optimization problem in both \bar{f} and σ_f may be solved using a suitable iteration method (see e.g. appendix C). We assumed here that the pulsar is sufficiently described by a log-normal distribution and that the dataset converges to the same (global) fix point for all initial pairs of values.

2.3 Detuning model

The problem of determining ToAs is closely related to a phase shift over the analyzed single pulses. As a wrongly set folding period can cause phase-drifts over the dataset's pulses, we will study the influence of such a detuning on the coefficients of a signal to exclude mistaking a phase-drift over the dataset as a phase shift when determining ToAs. Analysing Fourier transformations of finite (consecutive) pulses allow us to measure a phase-drift corresponding to the detuning of the assumed periodicity τ and the real periodicity $T = \frac{\tau}{a}$, where a is a correction factor. The equations denoted with a letter before the number are derived in appendix A2. For the signal from a single Fourier coefficient with frequency $\omega' = \frac{2\pi n}{T} = \frac{2\pi n a}{\tau}$ we will get a Fourier coefficient over the period $[K\tau - \frac{\tau}{2} : K\tau + \frac{\tau}{2}]$ of pulse K :

$$\tilde{d}_{k,K} = s_n \exp[2\pi i k (\frac{n}{k} a - 1) K] \cdot \text{sinc}[\pi k (\frac{n}{k} a - 1)] \quad (18)$$

where s and d with subscripts k, K denotes the k^{th} complex Fourier coefficient of the K^{th} sample. For $n = k$ eq (18) reads

$$\tilde{d}_{k,K} = s_k \exp[2\pi i k (a - 1) K] \cdot \text{sinc}[\pi k (a - 1)] \quad (19)$$

which means that, since $a \approx 1$, the sinc-factor is approximately 1 for $k \ll \frac{1}{(a-1)}$. Thus the relative phase of the K^{th} and the $(K + 1)^{\text{th}}$ pulse gives access to the model parameter a .

Let us now derive the Fourier coefficients for the full spectrum of a periodic, purely real signal. for such a signal, $s_n = s_{-n}^*$. Inserting

this relation and denoting the phase as $\arg(s_n) =: \Phi_{s_n}$ we get:

$$d_{k,K} = \sum_{n=1}^{\infty} a_n \text{sinc}[\pi(na - k)] \cdot \left[\underbrace{\exp[i(2\pi naK + \Phi_{s_n})]}_I + \underbrace{\frac{d}{k} \exp[-i(2\pi naK + \Phi_{s_n})]}_{II} \right] \quad (20)$$

where $\bar{k} = \frac{an+k}{2}$, $d = \bar{k} - k$. This formula can be written as $d_{k,K} = \sum_{n=1}^{\infty} A_{kn} a_n$. The main phase information again is to be found in term I of (21) for $n = k$. However this is not the only term relevant in analysing a single coefficient $d_{k,K}$. Assuming $n = k \pm m$ the sinc function slowly converges towards zero as $\frac{1}{m}$ and determines the factor for coefficient 1. Coefficient two together with the sinc function goes like $\frac{1}{2k \pm m}$ for $m > 0$. Thus in the limit of big difference in k and m , both terms equally contribute and the phase relation deforms more and more to an elliptic one.

We emphasize that for small detuning, the sinc-term is nearly zero for $k \neq n$. In this case we may calculate only with the diagonal matrix and thus correct the data for the phase $\Delta\Phi(a) = \exp[i2\pi naK]$. Thus for $a \cdot n_{max} \cdot K_{max} \ll 1$ the whole problem of detuning may be described, with a negligible error by the phase shift model in Sec. 2.6.

2.4 Phase model

Even for a perfectly tuned signal, we still have to face the noise introduced by the receiver and antenna system temperature on the phase. Again we will have to assign a measurement model for the data phase according to the signal. The noise passes the same linear transformation than the true Fourier coefficients and we conclude that

$$d_{k,K} = \sum_{n=1}^{\infty} A_{kn} (s_n + n_{n,K}) \quad (22)$$

where n is a white noise process with variance σ_n^2 . Since A is dependent on the true signal phases, the inverse problem as a whole is quite complex. But if we neglect term II of (21) we may invert the remaining matrix \tilde{A} without considering the true signal phases and apply it to a given dataset.

$$\underbrace{\sum_{k=1}^{\infty} \tilde{A}_{nk}^{-1} d_{k,K}}_{=: d'_{k,K}} = \|s_n\| \exp[i\Phi_{s_n}] + n_{n,K} \quad (23)$$

$$\tilde{A}_{kn} := \text{sinc}[\pi(na - k)] \exp[i2\pi naK] \quad (24)$$

This is possible since term II destroys phase information in a very smooth way and - as will be concluded below - high coefficient numbers will have higher phase errors due to the drop in signal amplitude and the relative growth of error. Consequently, higher coefficients do not give as much information about the phase and thus about a anyway. Having preprocessed the data with \tilde{A}^{-1} we arrive at the noisy signal and the phase carries information about Φ_{s_n} . Thus, given a and having processed the data vector, we may now fit for the average signal phase and variance by imposing a suitable model. We may estimate the error introduced in the true signal's phase by inspecting eq. (22) and Fig. 4:

$$\tan(\Delta\Phi) \approx \frac{\|n\|}{\|s\|} = \frac{\sigma_n}{\sigma_s \exp[2f]} \quad (25)$$




Figure 4. The phase error introduced to a signal by white noise scales as the expectancy value of the arctan of the signal to noise ratio

Rather than integrating this precisely one can use the following approximation without significant loss of accuracy

$$\sigma_{\Phi_s} := \langle \arctan(\exp[-2f]) \rangle_{\mathcal{P}(\mathbf{f}|\mathbf{d})} \quad (26)$$

$$\mathcal{P}(\Phi_{\mathbf{d}}|\Phi_{\mathbf{s}}) = \prod_i \sum_{k=-\infty}^{\infty} \exp\left[-\frac{(\Phi_{d'_i} - \Phi_{s_i} + 2\pi k)^2}{2\sigma_{\Phi_n}^2}\right] / C \quad (27)$$

This approximation is valid only for a good signal to noise ratio (S/N), where we also expect the most significant evidence to be. For low S/N it underestimates the error and we must numerically integrate the complete formula to get an error estimate. We emphasize that this is the very step where the confidence of the phase data and thus the weight in later calculations is calculated using the actual amplitude data point measured together with the statistics gathered. $\Phi_{\mathbf{s}}$ itself is parametrized by a so-called wrapped Gaussian⁴ and a mean of $\overline{\Phi_{\mathbf{s}}}$, and variance $\sigma_{\Phi_{\mathbf{s}}}$. we assume a prior of zero average and large variance.

2.5 Joint probability for phase

Finally we arrive at the joint probability of the phase model calculating analogue to the amplitude:

$$\mathcal{P}(\Phi_{\mathbf{d}}, \Phi_{\mathbf{s}}|T, \sigma_{\Phi_n}, a) = \exp\left[-\frac{1}{2} \left[\frac{\|(\Phi_{\mathbf{d}'}(\Phi_{\mathbf{d}}) - \Phi_{\mathbf{s}}) \bmod 2\pi\|^2}{\sigma_{\Phi_n}^2} + \frac{\|(\Phi_{\mathbf{s}} - \overline{\Phi_{\mathbf{s}}}) \bmod 2\pi\|^2}{\sigma_{\Phi_{\mathbf{s}}}^2} \right] \right] / C \quad (28)$$

Where $T := \{\bar{\mathbf{f}}, \sigma_{\mathbf{f}}, \sigma_{\Phi_{\mathbf{s}}}, \overline{\Phi_{\mathbf{s}}}\}$ will be the parameters of the full statistical template reconstructed and $\Phi_{\mathbf{d}'} = \Phi_{\mathbf{d}'}(\Phi_{\mathbf{d}})$ is the phase correction according to (24) or other corrections for the assumed phase shift. Integrating out $\Phi_{\mathbf{s}}$ we find:

$$\begin{aligned} \mathcal{P}(\Phi_{\mathbf{d}}|T, \sigma_{\Phi_n}, a) &= \\ &= \exp\left[-\frac{1}{2} \frac{\|(\Phi_{\mathbf{d}'}(\Phi_{\mathbf{d}}) - \overline{\Phi_{\mathbf{s}}}) \bmod 2\pi\|^2}{\sigma_{\Phi_n}^2 + \sigma_{\Phi_{\mathbf{s}}}^2}\right] / C \end{aligned} \quad (29)$$

⁴ a wrapped Gaussian is literally a gaussian distribution wrapped around the unit circle by identifying all points modulo 2π

We now again invert this relation using Bayes' theorem and take a MAP-Ansatz on our model parameters $\sigma_{\Phi_s}, \overline{\Phi_s}, a$.

2.6 Phase shift model for a single epoch as reference

In order to generate ToAs we will have to determine the relative shift in the time of arrival of the pulses compared to another epoch taken as a reference. For the sake of simplicity, we assumed that the pulsar period does not to change with time (which is only a matter of interpreting the relative phase correctly or further imposing the detuning model). The pulsar's signal is assumed to be slipped by a time Δt compared to the overall period. This amounts to a change in the phase of the Fourier coefficients given by:

$$d_k = \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} dt \quad d'_t \cdot \exp[-i\omega_k(t - \Delta t)] = d'_k \cdot \exp[i\omega_k \Delta t] \quad (30)$$

Thus, having reconstructed the template characteristics $T = \{\overline{f}_k, \sigma_k, \overline{\Phi}_k, \sigma_{\Phi_s}\}$ for the first time interval, the subsequent intervals may be analysed by carrying out a parameter study over Δt . This is done by shifting the measured phases $\overline{\Phi}_d \rightarrow \overline{\Phi}_{d'}$ according to (30) and taking e.g. a MAP-approach over the joint probability of the phase (29), now varying Δt , not a which is assumed to be 1. We now are also able to give the approximation for $a \approx 1$. We derived in this case a phase shift of $\Delta\Phi(a) = \exp[i2\pi naK]$. Comparing with (30), the problem of detuning for small a may be approximated by testing all pulses for a shift of

$$\Delta t = aK \cdot T \quad (31)$$

dependent on the pulsar period number K . All practical cases of detuning fall in the category of the aforementioned reduction to systematic phase shift.

2.7 Reference independent difference phase model

The simplest way to form a reference template is to select a single observation. However this relies on the chosen observation being statistically representative of the sample as a whole and would discard the extra information available in the whole dataset.

This in principle could be mitigated by building a total statistical template over all epochs observed. We discarded this method for two reasons. First of all it assumes that the pulsar template shape and statistical behaviour is the same in every single epoch (which is not the case for e.g. moding pulsars). Furthermore forcing the template to take the same probability density in signal amplitude for distinct epochs is too strict an assumption as for example ISM variations modulate the signal amplitude systematically. Instead, one may formulate a probability for the relative shift of two distinct epochs, given a certain epoch's statistical template is correct, and then demand that the probability distribution has to agree for all templates at once. We depicted this process schematically in Fig. 2.7. The calculated probability on a relative phase between two epochs, given all knowledge gathered, is a clear, basic statement coming without the notion of a "reference epoch". It can be incorporated in further analyses without introducing a direct bias caused by the phase information of statistical templates used since the exact phase assumed in the template drops out of the calculation. The knowledge gathered just filters the data to be compared in a clever way.

Mathematically speaking, we demand that

$$\begin{aligned} \mathcal{P}(\langle\Phi_d^{(i)}\rangle - \langle\Phi_d^{(ref)}\rangle = \Delta\Phi^{(i)}|T_1...T_n) &= \\ = \prod_j \mathcal{P}(\langle\Phi_d^{(i)}\rangle_{T_j} - \langle\Phi_d^{(ref)}\rangle_{T_j} = \Delta\Phi^{(i)}|T_j) & \quad (32) \end{aligned}$$

The formula for one factor of this product may be deduced from (29) by taking the product of the equation for two epochs' shifts $\Delta\Phi_1, \Delta\Phi_2 = \Delta\Phi_1 + \Delta\Phi$ and integrating out the difference to the used template, $\Delta\Phi_1$. A factor then reads:

$$\begin{aligned} \mathcal{P}(\langle\Phi_d^{(i)}\rangle_{T_j} - \langle\Phi_d^{(ref)}\rangle_{T_j} = \Delta\Phi^{(i)}|T_j) &= \prod_k \frac{1}{\sqrt{2\pi\sigma_k^{(i)2}}} \cdot \\ \cdot \exp\left[\frac{1}{2\sigma_k^{(i)2}}(\langle\Phi_d, k^{(i)}\rangle_{T_j} - \langle\Phi_d, k^{(ref)}\rangle_{T_j} - k\Delta\Phi)^2\right] & \quad (33) \end{aligned}$$

Where $\sigma^{(i)2} = \sigma_{s',k}^{(i)2} + \sigma_{s',k}^{(ref)2}$ is the sum of the new phase variances estimated from the epochs' data, given T_j holds. The formula deduced simply states that the difference of the filtered mean phases of the two epochs under consideration is the mean phase difference times the Fourier coefficient number.

While the gathered probabilities on phase differences are reference free, in order to output a TOA, we have to define that the absolute phase of a certain epoch amounts to zero. In the classical procedure, the template was assumed to have zero phase. Additionally there exist several conventions on defining the physical point on a profile where the phase is zero, including but not limited to defining the mid of the profile at the "center of mass" of the profile or at the highest peak. To that extent we declare a "reference epoch" which's ToA is declared to be exactly the timestamp of the rising edge of the first bin of the pulse. As we assume the epoch with the highest S/N-ratio to have the lowest variances in its statistical template as seen from the others, we choose this one to be the reference epoch. While, per definitionem, this TOA's relative phase is exactly zero, the corresponding ToA has to have a variance to be fitted into TEMPO2. Thus we calculate the variance of the reference as it would have been determined by all other epoch's relative shifts as

$$\frac{1}{\sigma^{(ref)2}} = \sum_j \frac{1}{\sigma^{(j)2}} \quad (34)$$

For the other epochs, the probability distribution of the relative phase shift provides the correct error estimate on the ToA automatically.

Another big advantage of parametrizing the ToAs by their difference is that (32) can be calculated independently for every template and added logarithmic. This can be done since the measured variable, phase difference, is the same for every pair of epochs, no matter which statistical template is assumed for the moment. This gives us the freedom to rasterize the probability distribution as a whole without preferring one template over the others and stay in the Bayesian picture until reducing the probability distribution measured into mean phase differences and variances in the very last step. This becomes essential when also short epochs with weak signal over noise ratio are to be included in the analysis. In this case multiple phases are still plausible e.g. mistaking one peak for the other (see also fig. 7 and Sec. 3.2) since the yielded probability distribution for one template may be too far from the gaussian form to be reduced to a mean and variance. The ToA t_{ToA} of epoch i with mean phase shift $\Delta\Phi^{(i)}$ and variance $\sigma^{(i)2}$ w.r.t. the declared




Figure 5. Overview of ToA generation: The statistical imprint of the various epochs (A) is used to generate statistical templates (B). The data of the epochs (D) is non-linearly filtered using the statistical templates to yield a probability distribution (C) for the shift of every epoch relative to a reference epoch. For every epoch, the mean phaseshift in topocentric pulsar periods is then added to epochs MJD to calculate the specific ToA (E). The error is calculated as the variance of (C) for this shift. The reference epoch's error on the ToA is determined as if it would have been timed by every other epoch.

reference then amounts to:

$$t_{\text{ToA}} = t_{\text{rising edge}} + P_{0,i} \times (\Delta\Phi^{(i)} \pm \sigma^{(i)}) \quad (35)$$

where $P_{0,i}$ is the folding period of epoch i and the phase difference takes values in the interval $[-.5 : .5]$.

Developing an interface to Bayesian extensions of TEMPO2, like TempoNest (Lentati et al. 2014), to communicate the whole probability distributions on the relative phase information to the pulsar timing code would be a further step to enhance the statistics on the ephemeris' parameter sets. While this may be desirable in the future we simply reduce the yielded probability distributions to their mean and variance and calculated ToAs directly comparable to the ones from tools like pat included in the PSRCHIVE toolset and processable by TEMPO2 without modifications.

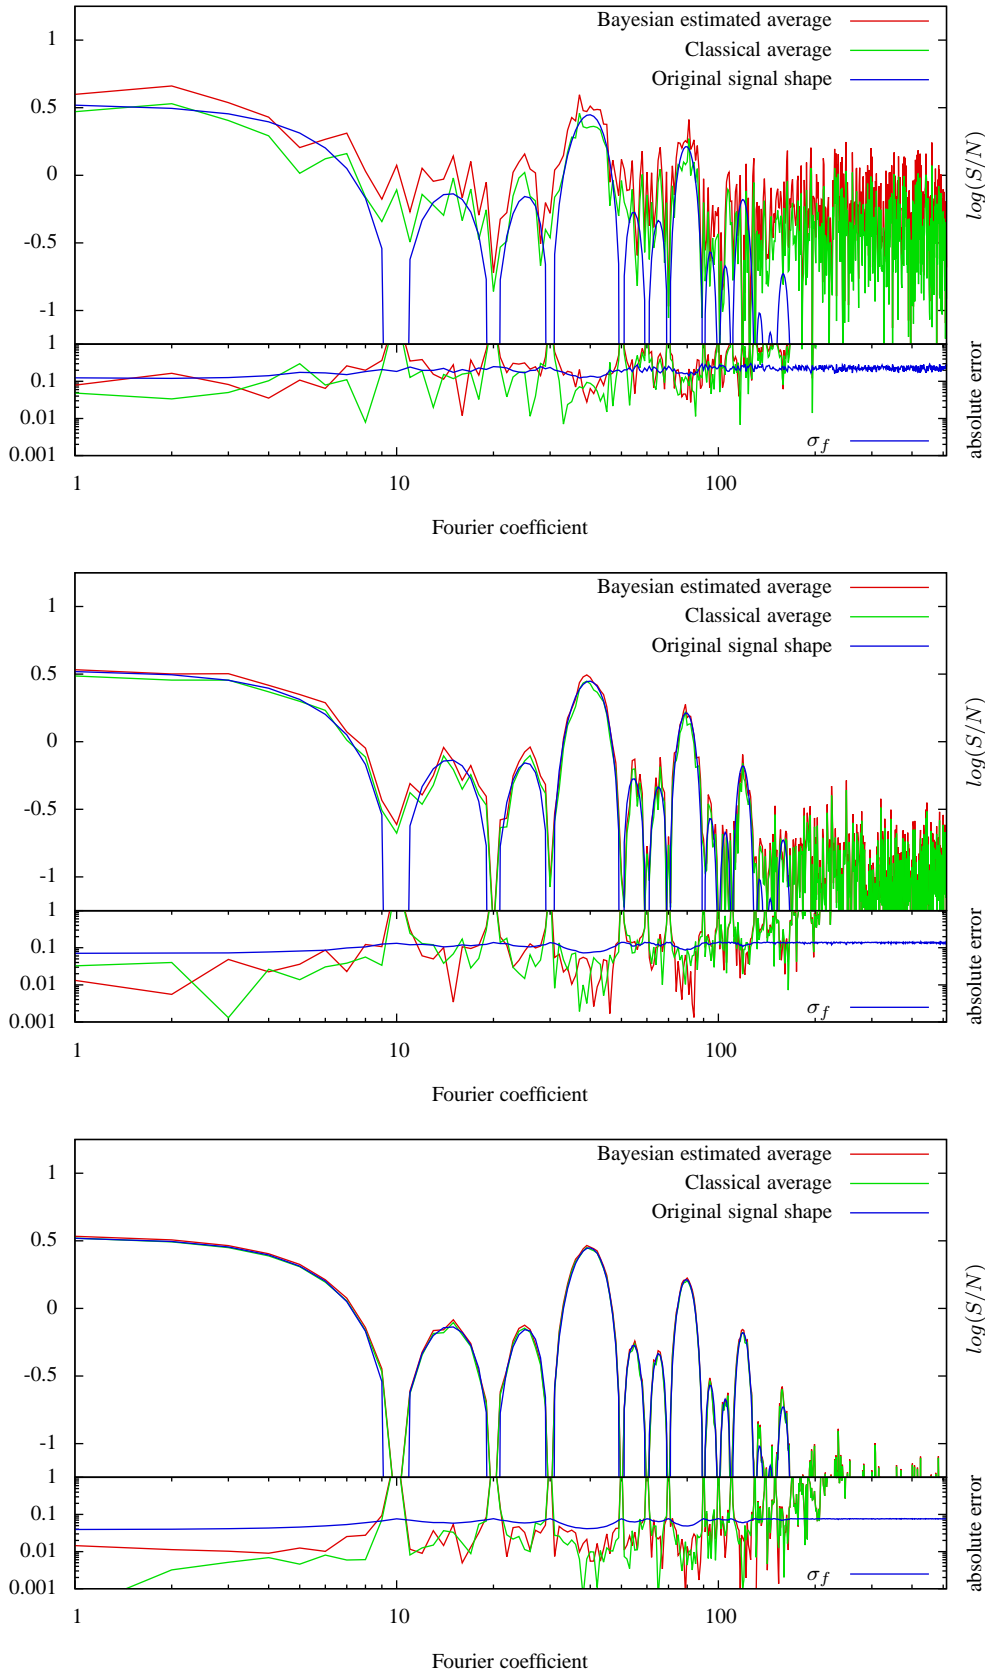
3 APPLICATION TO DATA

Having derived all relevant equations we may now apply the formulae upon given datasets and test the approach. Starting with simulated datasets we afterwards evaluate a dataset of PSR B1133+16 with 2041 consecutive single pulses consisting of 1024 bins each. Furthermore we determined ToAs from data of PSR J1713+0747

and analysed these ToAs with the TEMPO2 software package comparing with ToAs generated using “pat” from PSRCHIVE and a very accurate classical analytical template.

3.1 Simulated data

We wrote a simulated data generator serving single pulses of non-fluctuating shape \bar{f} but multiplied with gaussian red noise to account for a highly fluctuating radiation process and analysed the convergence of the amplitude model in Fourier space. The classical average is in this case the statistically optimal procedure to determine the pulse profile since the assumption that there is exactly one profile holds, and a quickly converging law of large numbers applies. Fig.6 depicts the reconstruction from ten, hundred and a thousand pulses comparing classical template integration with Bayesian reconstruction. Both methods perform equally well, as was expected. However, the Bayesian reconstruction also gives us an estimate on the profile stability and thus a better understanding of the remaining statistical uncertainty. Reconstructing a constant signal shape is a difficult task for a log-normal Bayesian model since the guess of a non-fluctuating shape is statistically only feasible by collecting lots of data. Thus, even though the test seems



After ten pulses, the Bayesian method has equally well reconstructed the signal as the average method, however it provides a handle on the signal fluctuations still possible. The seemingly systematic overestimation of the signal shape by the expected mean amplitudes comes from the log-normal exponential model used and the non vanishing σ_f

A hundred pulses already give a profile accurate to over 90% over the essential Fourier modes of the signal. Notice how the real error on the averaged \bar{f} is limited to the still possible statistical fluctuations σ_f .

A thousand pulses give a confident template up to high coefficients. The Bayesian reconstruction estimates the average shape to fluctuate below 20% of its mean. The absolute error not detected at the very high coefficients may be explained due to low assumed signal and low real signal to noise ratio. The dominant error source seems to be a systematic overestimation due to the log-normal signal model.

Figure 6. Reconstructions of the original signal shape by both Bayesian and classical methods and their errors after 10, 100 and 1000 pulses respectively. The lower half shows the errors of the estimation to the original signal shape and the Bayesian estimation of the remaining uncertainty σ_f on the expected log-normal mean \bar{f}

simple, it is actually a very strict test for the Bayesian measurement model to pass as the correct solution for infinite observation time is delta peaked. The algorithm is expected to perform much better on real data for which the profile shape measured is much more fluctuating for small integrations since it has a way of incorporating the instabilities of the profile. Since in an observation of finite length the algorithm will never show zero variance on the profile shape, it estimates the profile intensity slightly but systematically too high. This gives rise to a problem if one is set out to measure the intensity of radiation naively by interpreting the statistical template as a way of describing the profile. The template gives a statistical way to determine all compatible intensities or the probability distribution of these, but not the “true” intensity directly.

The high modes exhibit a large absolute error. This is an effect of low signal to noise ratio at imperfect numerical integration. For low amplitudes, the Hamiltonian for different values of \bar{f} becomes rather flat and thus the outcome just reflects the prior we have been taken. In this case the prior was a log-normal Gaussian model with $2\sigma_f$ around the average found. Notice that the σ_f the algorithm suggests (leaving aside the singularities of the profile) captures the true error (except for singularities and the imperfect numerical integration of high modes) giving a handle on the remaining statistical uncertainty.

3.2 Determining phase shift

We used a template with four peaks of Gaussian shape, sizes ranging from 1.8° to 3.6° and generated a dataset with an overall signal over noise ratio of $-17\text{dB} \approx \frac{1}{50}$. We compared two independently generated datasets of one hundred pulses with 1024 bins each, shifted by 1.8° , which amounts to a shift of 5.12 bins. As Fig. 7 shows, the Bayesian reconstruction yields $\Delta\Phi = -1.81^\circ \pm 0.04^\circ$. The other maxima in the probability distribution are of smaller size but reflect the the spacing of the templates peaks and the possibility of mistaking one peak for another. Having a local signal to noise ratio of four for every peak and a resolution of $360^\circ/1024 \approx 0.35^\circ$ available, we may estimate the uncertainty of every peaks location to be 1.3^{55} in every single peak in every single pulse. Having derived the phase drift using 200×4 peaks, we classically calculate the remaining error due to the receiver noise to $\frac{1.3^\circ}{\sqrt{2 \cdot 400}} \approx 0.046^\circ$. In the fluctuation free case (the simulated emitted signal is exactly the four peaked shape before the receiver noise is added and the signal is squared) it can be shown that this error is the minimum reachable error. The test of our algorithm reaches this accuracy, too.

3.3 Real data

Generating templates from real data yields reconstructions of surprising confidence. Fig. 8 depicts the amplitude reconstructions from certain numbers of consecutive pulses. In each figure, the long-term convergence gathered from the whole dataset may serve as a reference. We emphasize that for data gathered by a pulsar with fluctuating shape, classical and Bayesian reconstruction are in general not expected to yield the same curve, since fluctuations are

⁵⁵ Calculating the exact timing precision achievable is highly nontrivial. We estimated this error as follows: Assume a single peak at a S/N of one to be locateable to about its width, having a signal to noise ratio of four should increase this precision by a factor of $\sqrt{4}$. Adding the errors on the bin and resolution quadratically, we end up with an average of 1.3° per peak.





Figure 7. Unnormalized logarithmic probabilities for phase shifts. The comb structure comes from the four peaks present in the template.

handled differently by both models. Both can only be compared to their own long-term limit. Furthermore the statistics of the profile seem to change with time for the dataset examined. We will analyse this short term behaviour in the next section.

The Bayesian reconstruction does exhibit lower noise at high fourier coefficients and converges rather quickly to its ultimate form. The main features such as peaks and curvature appear already at 64 pulses while the classical average remains dominated by fluctuations occurring. However systematic offsets in the Bayesian reconstruction appear. These may largely be backtracked to the actual pulsars intensity changes and nulling periods. A numerical fluctuation of the algorithm is rendered unlikely since each point of the construction is calculated independently and thus an equal rise in all coefficients is highly improbable. The classical average also shifts upwards and downwards with varying intensity, but the effect is mostly indistinguishable from noise. For this real life example, the method derived is clearly superior to the classical method in areas of low signal to noise ratio, given that it converges very fast to its ultimate form and showing denoised systematic profile changes. It is reproducing the classical outcomes in the lower coefficients and not picking up the noise of the higher coefficients. After 128 pulses a confidence of about ten percent is already reached for all coefficients present. This paves the way to investigate changes in the radiation profiles on these scales. We could further diversify our pulses in intensity classes and perhaps in different modes of radiation given the contrast and fastness of the envelope estimate. Fig. 9

**Figure 8.** Template reconstruction for a subset of pulses

shows the differences of the two reconstruction methods in time domain and the single pulse probability density function as described by the Bayesian statistical template (depicted as grey shade): While the classical average just assumes the pulsar to have a certain profile shape, the Bayesian method is rather based on exclusion and

describes profile shapes compatible with the data gathered. Thus, after as few as ten pulses, the Bayesian reconstruction in time domain already resembles the final form of the radiation profile. Notice that it still allows for signals out of the profile, since statistics do not exclude low signals after ten pulses.




Figure 9. Template reconstruction after ten respectively. a thousand pulses. While the classical averaging assumes a certain fluctuating form to be true, the Bayesian template and single pulse model shows which single pulses are typically compatible with the statistics of the data gathered




Figure 10. Depicted are the Bayesian and classic reconstructions from two sets of 50 pulses shifted by a single pulse. The change of the classical template is dominated by the spiky fluctuations as is evident from the change in the right peak.

3.4 Reconstruction stability to fluctuations

The workings and main differences of the two algorithms become evident when analysing the change of the classical profile and statistical template introduced by small changes of the dataset. This is examined by comparing the reconstructions of a window of fifty pulses moved along the dataset. An example of stepping one pulse

further is shown in Fig. 10. Dropping one pulse and inserting another should not change the overall template much. If the template changes, we expect it to change smoothly with time. For the Bayesian reconstruction, this property can be easily observed whilst individual fluctuations dominate the classical average picture. This can be deduced from the right peak of the template. While the algorithm derived gives a rather smooth increase in intensity

over the whole peak, the classical average is dominated by spikes of individual fluctuation. One could argue that the classical spikes might just be smoothed out of the form of the Fourier transform. This is contradicted by highly localized changes in the Bayesian reconstruction that are statistically significant. The change of the left peak may be seen as an example for a localized statistical change. The classical average does not change much there, but the Bayesian reconstruction shows a significant rise of the left flank of the peak. The algorithm is found to behave in such a way over the whole dataset. Furthermore it incorporates the fluctuation seen in real envelopes. This may increase accuracy of derived parameters like the ToAs since e.g. the confidence of a peak appearing at a certain point should be, but is not, taken into consideration by classical timing procedures.

3.5 Nulling

Given a set of proposed models, like different statistical templates, Bayesian statistics can also assign a probability on how likely every single model is to describe the data, given that one of the proposed models is true. An example for the ability to discriminate between different models is nulling detection. Two sets of parameters are analysed within the framework: One set of \bar{f}, σ_f is generated from a training set of nulling periods taken from the first 50 pulses, the other one is taken from the overall reconstruction as above. The likelihood of a model i described by a set of parameters m_i is compared as

$$\mathcal{P}(m_i|\mathbf{d}) = \frac{\iint \mathcal{D}\mathbf{s} \mathcal{P}(m_i, \mathbf{s}, \mathbf{d})}{\sum_i \iint \mathcal{D}\mathbf{s} \mathcal{P}(m_i, \mathbf{s}, \mathbf{d})} \quad (36)$$

which may be derived straight-forwardly from Bayes' theorem. Deciding between the white noise model and the model derived from all thousand pulses, this formula gives the probabilities given in Fig. 11 for a certain frame to be a nulling pulsar period. These probabilities could be used as weights when determining further parameters without biasing the statistics by deciding whether there was a pulse or nulling. Furthermore when looking at the plot, the Bayesian probabilities resemble our state of uncertainty for pulses which are in the 50% region when a by eye decision should be made. Fortunately, the method is able to quantify this uncertainty for us automatically. Nulling information could in principle further improve the accuracy of detuning and time of arrival analysis.

3.6 Determining Times Of Arrival

3.6.1 ToAs from simulated data

Timing simulated data sets is a precise way to test the algorithm for theoretical performance since the true signal to be found out of noisy data are known. We implemented the phase shift model (as described in Sec. 2.6). We generated a statistical template out of 10,000 simulated pulses and then measured the accuracy on timing $N = 10, \dots, 50,000$ pulses with that fixed reference, where we used a MAP approach. Statistics was gathered over 20 randomly shifted datasets for each value of N recording the absolute error on the ToA and deriving the mean error and standard deviation to quantify the accuracy of the algorithm. The algorithm performed as depicted in Fig. 12. As expected, the systematic phase error was well bounded by the fluctuations given by the test datasets. The deviation from the true value followed a $1/\sqrt{N}$ law for large N as expected and reached the accuracy that is information theoretical possible. When there were more than 10,000 pulses in the dataset

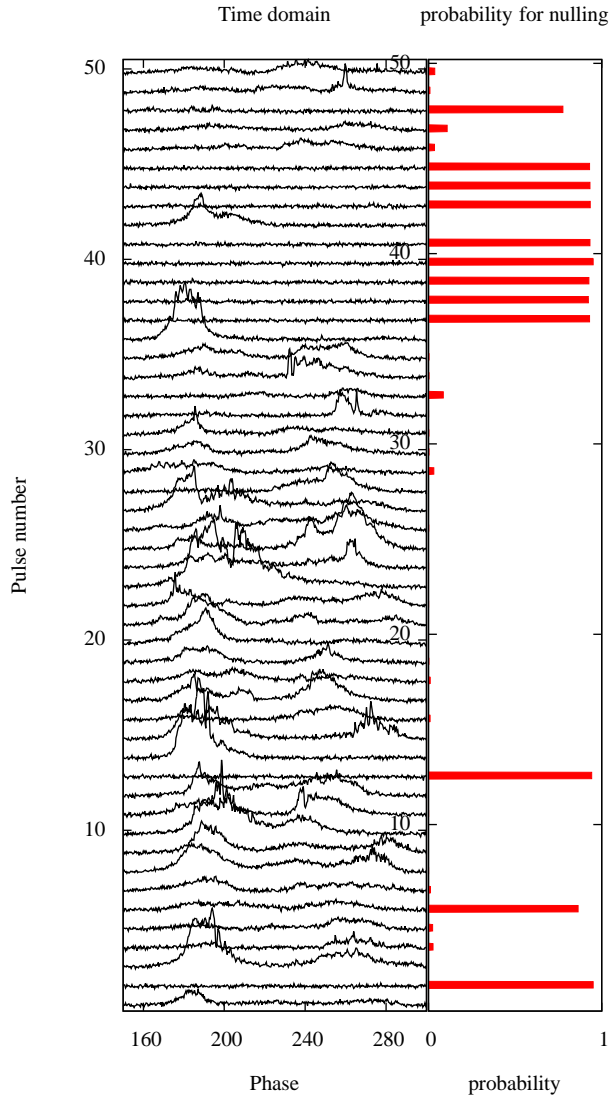


Figure 11. Discrimination-less decision: Bayesian analysis gives a probability for a certain pulse candidate to be explained by white noise only. In subsequent data-processing, this probability may be used to weight signals that are only present in non-nulling phase instead of a threshold algorithm producing false positives or negatives.

to time, a systematic error appeared. This was expected, since the reference statistical template itself (generated from 10,000 pulses) does not contain more accurate phase information. Since we were measuring an absolute, not a relative shift here, the error contained in the template becomes systematic. The sudden drop of the systematic error between 10 and hundred pulses is a random sampling artifact of the simulation involved. The average over the discrete maximum values taken was in this case accidentally zero. We also evaluate the reached accuracy for different S/N-ratios in Fig. 13. It is interesting to notice that the algorithm returns a rather flat probability distribution rendering a lot of phase values likely. This can be deduced from the high errors encountered leading to a plateau for low values of N . At a certain point the probability distribution begins to peak and accuracy increases with a jump, locking onto the signal. This may be explained by passing the point where different peaks of the profile shape can be distinguished statistically. This is




Figure 12. Test of accuracy over epoch size of the phase shift model as described in sec. 2.6. a variable number of pulses is used to generate ToAs and the algorithm is examined for statistical (red) and systematical (green) errors.




Figure 13. Test of accuracy for different S/N-ratios of the phase shift model as described in sec.2.6 the grey dashed lines all follow a $1/\sqrt{N}$ slope.

not a drawback of the method but reflects the fact that, also in the ideal case, the error follows the $1/\sqrt{N}$ slope only for large N and a unique maximum. After that point the performance follows a stable $1/\sqrt{N}$ curve.

We also simulated the whole process from measuring epochs to determine the ToAs therein. To this end we simulated 20 epochs with 100 pulses each. Then we generated statistical templates from each epoch and used that to time the relative shifts of every other epoch. Then we naively averaged the maximum a-posteriori values of the shifts and their variances using classical formulae for




Figure 14. The upper panel depicts the remaining uncertainty (in an arbitrary scale) when the row epoch is measured by the column epoch. Notice that there are some combinations of epochs which lack statistical similarities and give greater uncertainties. The lower plot shows the remaining error on the ToAs of the epoch when combining the measurements from above.

measurement value addition to find the ToAs and compare them with their true values. The dataset contained 2000 single pulses. The results of this simulated measurement is displayed in Fig. 14. We reach an accuracy of 1.9×10^{-3} radians, which is nearly the theoretical maximum for 100 pulses per epoch as can be seen from Fig.12. When timing with only 10 epochs, the accuracy dropped to about 3×10^{-2} radians. In that case, the theoretically possible accuracy for each epoch was missed. That practically means, that gathering knowledge of later epochs can improve measurements of the past by reprocessing them with the newer statistical templates.

3.6.2 ToAs from real data

The method carried out in such a way fails when tested with real data. Reducing the phase shift information of every statistical template to classical mean shifts and their variances before correlating the different templates' measurement flaws the precision of the method. This can be understood easily considering very noisy templates. These may cause different very likely shift values separated by a larger interval of unlikely phase values, e.g. different peaks of the profile may be mistaken for each other. When reducing, the mean and variance of only the most likely peak was determined giving a rather low sigma value on it, since in its neighbourhood, the probability density can be approximated as

Table 1. Comparison of timing residuals of observations taken at 1350MHz and model accuracy as reported by TEMPO2 for classical and Bayesian ToAs. The smallest values in each category have been highlighted.

TEMPO2 fitting mode	rms/Wrms	χ^2	Classical timing		Bayesian timing	
			red. χ^2	rms/Wrms	χ^2	red. χ^2
RMS	551ns	-	-	431ns	-	-
weighted RMS	240ns	455	23	163ns	35	2

a gaussian. This sets a too large weight on the statistical template compared with the other templates if the algorithm accidentally locked on the wrong peak. The consequence is a large bias in the direction of the mistaken peak. However, if one does not reduce the data template by template, but examines the probability density on the relative phase shift of two epochs using all statistical templates at once, one circumvents the possibility of this local fallacy and gets an overall correct probability distribution.

An algorithm obeying these caveats has been outlined in Sec. 2.7. By measuring the same quantity (relative shift of two epochs) with different statistical templates we may rasterize the probability distribution for that quantity considering all references at once and in parallel. Since pulsar timing expects us to state a single ToA per epoch and the expected error on it, the MAP ansatz taken in the very last step is now justified. The reported errors on the ToAs are found to be in agreement with the RMS reported by TEMPO2 yielding low χ^2 values.


Using 10 second observations of PSR J1713+0747 at 1350MHz with a total of about 19h observation time in 36 epochs observed in Effelsberg over about 1.9 years we tested the algorithm. Even though single pulse observations carry much more statistical information than the data at hand, the statistical templates generated show variations. The ToAs generated reach an RMS of 431ns and a weighted RMS of 163ns at a fit χ^2 of 35. The classical method yields 551ns unweighted or 240ns weighted RMS with a χ^2 of 455. We collected the results in Table 1.

3.7 Moding

We examined the algorithm’s timing behaviour on a moding pulsar. We used a dataset of pulsar B0329+54 at 21 cm consisting of 5000 single pulse observations with 1024 bins each. Using the same formula (36) as for the nulling analysis, the algorithm decides if a single pulse or a subset of pulses was emitted in a certain pulsar mode described by statistical template m_i (consisting of the same set of parameters like T_i). We auto-generated a moding analysis by generating a statistical template T over all data. Then we analysed the probability of each given single pulse to appear, given the template T describes the radiation correctly. Basically this means evaluating the joint probability for the phase (28) and amplitude (by integrating (11) over s). We ordered this list of probabilities descending and initially divided it into n_{modes} sections of equal length, where we assumed a certain n_{modes} to be the correct number.

In principle, the number of statistically distinct modes is also subject to uncertainty. One may also calculate the probability, that a certain n_{modes} holds. This could in principle further improve the results derived in the following but is of secondary interest for a first analysis of the behaviour of the algorithm.

Given this initial assignment of the single pulses to modes, we started an iteration procedure. In each step, a new set of statistical templates according to the weights of the previous step are calcu-

**Figure 15.** We overlayed the average profile curves of different modes over a scatter plot of the single pulses coloured according to the probability of belonging to mode 1-3. The dotted shape amounts to an average over all pulses.

lated: $\{T_1^{(0)}, \dots, T_{n_{modes}}^{(0)}\}$. These were then used as models in eq. (36) to assign the probability for each individual pulse to belong to a certain mode. These probabilities become the new weights for that pulse. We experimented with this iterating procedure and found about five steps to be sufficient that the difference from step to step is negligible.

The algorithm allows an unambiguous assignment of single pulses to a certain statistical template except for a very small number of pulses. The statistical nature of single pulses makes it difficult to grasp what classification is happening. Thus we decided to make a scatter plot (Fig. 15) of 5000 single pulses assigned to three modes. The single pulses were blurred to simulate a density kernel and plotted with an opacity of 1% in the colour of the probability to belong to a certain mode. For orientation, we overlayed the average pulse shapes each mode would generate if it was the sole mode observed along with an average over all pulses, drawn in dotted white.

The term moding usually refers to the appearance of few distinguishable shapes of the integrated pulse profile. As a single pulse profile is very different from its neighbouring pulses, it was very common to assign different modes to consecutive pulses. Thus the question arises, whether to still call this behaviour moding or not. On one hand, we could try to further develop the algorithm to make mode switching on such small timescales very unlikely. Then the algorithm would pick up “modes” in a more classical sense. On the other hand, the integrated profiles of these “sub”-modes are quite distinct and certainly their ratio has an impact on the average profile

25 vs. 50 pulse integration window




Figure 16. This figure compares the phase detection performance of a single or a moding statistical template on 25 respectively 50 pulse integrations.

shape over a few minutes integrating and thus also on timing (as we will examine below). Understanding the relationship between these sub-modes and the astonishingly stable moding behaviour over a larger integration time could lead to a deeper understanding of the conditions in the pulsar magnetosphere.

These submodes also seem to be correlated over the whole profile. For example, mode three in the figure is the only one having an earlier rise of intensity at 230 degrees and additionally a very low slope in the middle of the profile.

Taking these different correlations into account can also archive a more accurate timing and make the timing stable against different integration times and noise. Fig. 16 shows two examples of how a moding template on single pulse level can improve detecting the phaseshift of a few tens to hundreds of pulses. Statistical templates were generated from 5000 pulses and later used to time simulated epochs consisting of consecutive subsets of 25 in the upper and 50 single pulses in the lower panel, taken from the same 5000 pulses. The most likely value for the shift of each set was binned to generate histograms of the observed frequency. Knowing that there was no subpulse drift over the 5000 pulses as a whole, we expect the phase to be measured as zero for every subset. This is indeed the case, if we assume the pulsar to radiate in different modes (depicted in green). When run over subsets of 25 pulses, the non-moding template (depicted in red) fails to detect the correct shift by repeatedly mistaking one peak for the other. Unfortunately the few outliers which would have identified the right peak are scattered around zero with a systematic bias. We suspect this bias to be caused by the single statistical template trying to cover two or more very distinct modes of radiation. Doubling the number of pulses in a subset fixes the problem of not matching the right peaks. However, there is still a bias and the results have a larger variance than the moding template. The moding template detects the correct pulse phase with satisfying accuracy while showing a lower variance. This can be understood having a look at the generated modes. The maximum value of the second peak for every single mode is further left for pulses which are less intense. In the average picture over all pulses this correlation cannot be accounted for. When looking at subsets

of pulses however, the lower probability of reaching a high intensity mode shifts the peak to the left. Consequently, the reference template is detected to be shifted to the right. The ansatz using one statistical template reduces this inaccuracy to about half a degree, which is still a low value in the light that e.g. the peaks of the first and second mode, if we assume three modes (see Fig. 15), are separated by 1.5 degrees and the average taken from 50 pulses is still very noisy. Using a statistical moding template, the variance in variability exceeds the inaccuracy introduced by integrating a smaller subset than the one the template was generated from.

We conclude that assuming even a low number of modes to be present can significantly improve the timing results both in variance and systematic error.

4 CONCLUSIONS

We developed and evaluated a log-normal Bayesian model for single pulse analysis of pulsars. The algorithm described is able to reproduce the results of classical averaging procedures when considering non-fluctuating templates up to the approximations taken. The method is surprisingly versatile in situations of weak signal or a fluctuating template shape, since it is shown to be more robust against fluctuations. Thus it may enable one to examine fast changes in the radiation profile or distinguish different modes of radiation, such as the nulling analysis, an example for model discrimination, has shown.

Secondary parameters' reconstruction may be implemented easily by calculating their imprint on the data and then evaluating the joint probability using the model derived, as was demonstrated with phase drift on simulated epochs. The parameter studies derived in such a way profit from additional statistical information incorporated in the template automatically since only the statistically significant part of the data will have an imprint on the probability distribution of the parameters under examination.

The benefit of rather understanding which parameter values are compatible with the data gathered than trying to reconstruct one possible signal out of the dataset might be worth the effort im-

plementing a Bayesian reconstruction for the parameter of your choice. This may give statistical access to hitherto poorly examined parameters and topics about the radiation process and magnetosphere by providing strong and reasonable reconstructions for noisy and fluctuating data where a classical estimator leaves us with a definite but perhaps insignificant answer.

A first test to generate ToAs from real data showed that the algorithm developed can reduce χ^2 values while reaching the accuracy of classical methods even though we did not use single pulses but few second integrations and approximated some calculations. It gave correct ToAs even though the statistical templates used were generated from the same data that was to time. The classical way of template generation is known to affect the results of timing if the same data is used for template generation (Hotan et al. 2005).

If one intends to use the method on routine basis as an alternative to the classical one, the data volume to keep for logging and generating reproducible output increases drastically. One would have to find a way of reducing the data without losing the possibility to reprocess it or re-validate it.

Data compression and performance might be improved at the same time for Bayesian methods, as van Haasteren (2013) already points out. Computational feasibility sets a limit to the potential use of Bayesian analysis, but there exist ready to use methods for speeding up Bayesian calculations, such as Taylor et al. (2012) or the Metropolis-algorithm (Metropolis et al. 1953). The formulae presented may also be sped up by analysing them in perturbative manner. Enßlin et al. (2009) and references therein develop an information field theory, taking the information Hamiltonian $H = -\log \mathcal{P}(\dots)$ as a starting point. This allows for approximating the numerical integrations involved analytically and breaking their influence down to computationally simpler matrix arithmetic.

Furthermore consolidating different data channels, like radio data at different frequencies, is easily possible in a Bayesian framework. Once the computational complexity is dealt with, the Bayesian method might provide a way to reduce observation time for low signal to noise pulsars and give a handle on the statistical uncertainties involved in template generation. It gives access to a broad range of secondary applications such as sophisticated methods of interference detection or testing pulsars long-term behaviour for template changes or short- and long-term moding behaviour affecting ToA-analyses.

Hence the basic analysis carried out has shown that Bayesian reconstruction of pulsar templates is not only feasible, but also forms a more complete picture on the pulsars template shape. The additional data turns out to be very valuable for subsequent analysis of secondary parameters mitigating the effects of insignificant fluctuations. Benefits known to arise by inspection of the stochastic behaviour in the analysis will be automatically propagated by implementing a measurement using the proposed Bayesian template. Having these manifold possibilities in mind we suspect the method to bear fruit worth picking and suggest further examination.

5 ACKNOWLEDGEMENTS

The authors would like to thank Thorsten Enßlin for valuable input on the measurement model. We thank Axel Jessner to provide a dataset of B1133+16 conveniently preprocessed to access in the code and valuable discussion. We thank Kejia Lee for the extensive comments shaping and sharpening the argumentative line of the paper. We thank Niels Oppermann for extended discussion on the

mathematical difficulties associated with the evaluation of the posterior pdf.

REFERENCES

- Bailes M., 2010, in Klioner S. A., Seidelmann P. K., Soffel M. H., eds, IAU Symposium Vol. 261, IAU Symposium. p. 212–217
- Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, MNRAS, 418, 561–570
- Cordes J. M., Kramer M., Lazio T. J. W., Stappers B. W., Backer D. C., Johnston S., 2004, New A Rev., 48, 1413–1438
- Cordes J. M., Shannon R. M., 2010, ArXiv e-prints
- Dempster A. P., Laird N. M., Rubin D. B., 1977, Journal of the Royal Statistical Society. Series B (Methodological), 39, pp. 1–38
- Edwards R. T., Hobbs G. B., Manchester R. N., 2006, MNRAS, 372, 1549–1574
- Enßlin T. A., Frommert M., Kitaura F. S., 2009, Phys. Rev. D, 80, 105005
- Helfand D. J., Manchester R. N., Taylor J. H., 1975, ApJ, 198, 661–670
- Hobbs G. B., Edwards R. T., Manchester R. N., 2006, MNRAS, 369, 655–672
- Hotan A. W., Bailes M., Ord S. M., 2005, MNRAS, 362, 1267–1272
- Hotan A. W., van Straten W., Manchester R. N., 2004, PASA, 21, 302–309
- Kramer M. et al., 2006, Science, 314, 97–102
- Kulkarni S. R., 1989, AJ, 98, 1112–1130
- Lee K. J. et al., 2014, MNRAS, 441, 2831–2844
- Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, MNRAS, 437, 3004–3023
- Liu K., Keane E. F., Lee K. J., Kramer M., Cordes J. M., Purver M. B., 2012, MNRAS, 420, 361–368
- Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, J. Chem. Phys., 21, 1087–1092
- Ng M., Pan J., 2010, SIAM Journal on Scientific Computing, 32, 1442–1464
- Ośłowski S., van Straten W., Hobbs G. B., Bailes M., Demorest P., 2011, MNRAS, 418, 1258–1271
- Rickett B. J., 1975, ApJ, 197, 185–191
- Shannon R. M. et al., 2014, MNRAS, 443, 1463–1481
- Taylor J. H., 1992, Royal Society of London Philosophical Transactions Series A, 341, 117–134
- Taylor S. R., Gair J. R., Lentati L., 2012, ArXiv e-prints
- van Haasteren R., 2013, MNRAS, 429, 55–62
- Vigeland S. J., Vallisneri M., 2013, ArXiv e-prints
- Wiener N., 1949, Extrapolation, Interpolation, and Smoothing of Stationary Time Series. Wiley

APPENDIX A: CALCULATIONS

A1 Integration over \mathbf{g}

Starting from eq. (10):

$$\begin{aligned}
H_B[A] &= \frac{1}{2} \left[(\sqrt{\mathbf{d}} - \widehat{\exp[\mathbf{f}]} \cdot \mathbf{g} \sigma_n)^\dagger \frac{1}{\sigma_n^2} (\sqrt{\mathbf{d}} - \widehat{\exp[\mathbf{f}]} \cdot \mathbf{g} \sigma_n) + \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f} + \mathbf{g}^\dagger \mathbf{G}^{-1} \mathbf{g} \right] \underbrace{- \log \mathcal{P}(\tau) + \frac{1}{2} \log(|2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |2\pi\mathbf{G}| |\mathbf{d}|)}_{=:\mathbf{A} - 0.5 \cdot \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f}} \\
&= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} \sqrt{\mathbf{d}} - 2 \frac{(\sqrt{\mathbf{d}} \widehat{\exp[\mathbf{f}]})^\dagger}{\sigma_n} \mathbf{g} + \mathbf{g}^\dagger \widehat{\exp[\mathbf{f}]}^\dagger \widehat{\exp[\mathbf{f}]} \mathbf{g} + \mathbf{g}^\dagger \mathbf{G}^{-1} \mathbf{g} \right] + \mathbf{A} = \\
&= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} \sqrt{\mathbf{d}} + \mathbf{g}^\dagger \underbrace{\left(\mathbf{G}^{-1} + \widehat{\exp[\mathbf{f}]}^\dagger \widehat{\exp[\mathbf{f}]} \right)}_{=: \mathbf{D}_f} \mathbf{g} - 2 \frac{(\sqrt{\mathbf{d}} \widehat{\exp[\mathbf{f}]})^\dagger}{\sigma_n} \mathbf{g} \right] + \mathbf{A} = \\
&= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} \sqrt{\mathbf{d}} + (\mathbf{g} - \dots)^\dagger \mathbf{D}_f (\mathbf{g} - \dots) - \left(\frac{\sqrt{\mathbf{d}} \widehat{\exp[\mathbf{f}]}}{\sigma_n} \right)^\dagger \mathbf{D}_f^{-1} \left(\frac{\sqrt{\mathbf{d}} \widehat{\exp[\mathbf{f}]}}{\sigma_n} \right) \right] + \mathbf{A}
\end{aligned}$$

Integrating over \mathbf{g} leads to:

$$\begin{aligned}
H_B[A \setminus \mathbf{g}] &= \frac{1}{2} \left[\frac{\sqrt{\mathbf{d}}^\dagger}{\sigma_n^2} \left(1 - \widehat{\exp[\mathbf{f}]}^\dagger \mathbf{D}_f^{-1} \widehat{\exp[\mathbf{f}]} \right) \frac{\sqrt{\mathbf{d}}}{\sigma_n} + \log |2\pi \mathbf{D}_f^{-1}| \right] + \mathbf{A} = \\
&= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} \left(1 - \tilde{\mathbf{D}}_f^{-1} \right) \sqrt{\mathbf{d}} + \log |2\pi \mathbf{D}_f^{-1}| \right] + \mathbf{A} = \\
&= \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} \left(1 - \tilde{\mathbf{D}}_f^{-1} \right) \sqrt{\mathbf{d}} + \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f} + \log(|2\pi \mathbf{D}_f^{-1}| |2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |2\pi\mathbf{G}| |\mathbf{d}|) \right] - \log \mathcal{P}(\tau)
\end{aligned} \tag{A1}$$

where $\tilde{\mathbf{D}}_f = \widehat{\exp[-\mathbf{f}]} \mathbf{G}^{-1} \widehat{\exp[-\mathbf{f}]} + 1$ is always greater than 1 leading to a broadening of the likelihood in case of bad signal noise ratio. We may reformulate the matrix in the data-dependent term as

$$1 - \tilde{\mathbf{D}}_f^{-1} = [\tilde{\mathbf{D}}_f^{-1}] [\tilde{\mathbf{D}}_f - 1] = [\widehat{\exp[-\mathbf{f}]} \mathbf{G}^{-1} \widehat{\exp[-\mathbf{f}]} + 1]^{-1} \widehat{\exp[-\mathbf{f}]} \mathbf{G}^{-1} \widehat{\exp[-\mathbf{f}]} = [1 + \widehat{\exp[\mathbf{f}]} \mathbf{G} \widehat{\exp[\mathbf{f}]}]^{-1} \tag{A2}$$

which leads us to the final form

$$H_B[A] = \frac{1}{2} \left[\sqrt{\mathbf{d}}^\dagger \frac{1}{\sigma_n^2} [1 + \widehat{\exp[\mathbf{f}]} \mathbf{G} \widehat{\exp[\mathbf{f}]}]^{-1} \sqrt{\mathbf{d}} + \mathbf{f}^\dagger \mathbf{F}^{-1} \mathbf{f} + \log(|2\pi \mathbf{D}_f^{-1}| |2\pi\sigma_n^2|^{N_{\text{tot}}} |2\pi\mathbf{F}| |2\pi\mathbf{G}| |\mathbf{d}|) \right] - \log \mathcal{P}(\tau) \tag{11}$$

A2 Discrete Fourier coefficients of a detuned signal

For the signal consisting of a single Fourier coefficient with frequency $\omega' = \frac{2\pi n}{T} = \frac{2\pi n a}{\tau}$ the measured coefficient over a period $[K\tau - \frac{\tau}{2} : K\tau + \frac{\tau}{2}]$ of pulse K is readily calculated as

$$\begin{aligned}
\tilde{d}_{k,K} &= \int_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} dt \quad s_n \exp[i(\omega'_m - \omega_k)t] = \\
&= s_n \left[\frac{\exp[2\pi i \frac{(na-k)}{\tau} t]}{2\pi i (na - k)} \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
&= s_n \left[\frac{\exp[2\pi i k \frac{(\frac{n}{k}a-1)}{\tau} t]}{2\pi i k (\frac{n}{k}a - 1)} \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
&= s_n \exp[2\pi i k (\frac{n}{k}a - 1)K] \frac{\sin \pi k (\frac{n}{k}a - 1)}{\pi k (\frac{n}{k}a - 1)} \\
&= s_n \exp[2\pi i k (\frac{n}{k}a - 1)K] \cdot \text{sinc}[\pi k (\frac{n}{k}a - 1)]
\end{aligned} \tag{18}$$

For the full spectrum of a real signal, where $s_n = s_{-n}^*$ we get

$$\begin{aligned}
 d_{k,K} &= \int_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} dt \sum_{n=-\infty}^{\infty} s_n \exp[i(\omega'_m - \omega_k)t] = \\
 &= \left[\sum_{n=-\infty}^{\infty} s_n \frac{\exp[2\pi i \frac{(na-k)t}{\tau}]}{2\pi i(na-k)} \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
 &= \sum_{n=1}^{\infty} \left[s_n \frac{\exp[2\pi i \frac{(na-k)t}{\tau}]}{2\pi i(na-k)} + s_n^* \frac{\exp[2\pi i \frac{(-na-k)t}{\tau}]}{2\pi i(-na-k)} \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
 &= \sum_{n=1}^{\infty} a_n \left[\frac{(k+na)(\cos \Phi_{s_n} + i \sin \Phi_{s_n}) \exp[2\pi i \frac{(na-k)t}{\tau}] + (k-na)(\cos \Phi_{s_n} - i \sin \Phi_{s_n}) \exp[2\pi i \frac{(-na-k)t}{\tau}]}{2\pi i(k^2 - (na)^2)} \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
 &= \sum_{n=1}^{\infty} a_n \left[\exp[2\pi i \frac{-k}{\tau} t] \left(\cos \Phi_{s_n} \frac{-ik \cos(\frac{2\pi na}{\tau} t) + na \sin(\frac{2\pi na}{\tau} t)}{\pi(k^2 - (na)^2)} + \sin \Phi_{s_n} \frac{ik \sin(\frac{2\pi na}{\tau} t) + na \cos(\frac{2\pi na}{\tau} t)}{\pi(k^2 - (na)^2)} \right) \right]_{K\tau - \frac{\tau}{2}}^{K\tau + \frac{\tau}{2}} = \\
 &= \pm \sum_{n=1}^{\infty} a_n \left[2 \sin(\pi na) \frac{\cos \Phi_{s_n} (ik \sin(2\pi na K) + na \cos(2\pi na K)) + \sin \Phi_{s_n} (-na \sin(2\pi na K) + ik \cos(2\pi na K))}{\pi(k^2 - (na)^2)} \right] = \\
 &= \pm \sum_{n=1}^{\infty} a_n \left[2 \sin(\pi na) \frac{i \sin(2\pi na K) (\cos \Phi_{s_n} k + i \sin \Phi_{s_n} na) + \cos(2\pi na K) (\cos \Phi_{s_n} na + i \sin \Phi_{s_n} k)}{\pi(k^2 - (na)^2)} \right] = \\
 &= \sum_{n=1}^{\infty} a_n \left[2 \text{sinc}[\pi(na-k)] \frac{i \sin(2\pi na K) (\cos \Phi_{s_n} k + i \sin \Phi_{s_n} na) + \cos(2\pi na K) (\cos \Phi_{s_n} na + i \sin \Phi_{s_n} k)}{k+na} \right] =
 \end{aligned}$$

where we now introduce $\bar{k} = \frac{an+k}{2}$, $d = \bar{k} - k$ and simplify

$$\begin{aligned}
 &= \sum_{n=1}^{\infty} a_n [\text{sinc}[\pi(na-k)] \cdot \\
 &\quad \cdot \frac{i \sin(2\pi na K) (\cos \Phi_{s_n} (\bar{k} - d) + i \sin \Phi_{s_n} (\bar{k} + d)) + \cos(2\pi na K) (\cos \Phi_{s_n} (\bar{k} + d) + i \sin \Phi_{s_n} (\bar{k} - d))}{\bar{k}}] = \\
 &= \sum_{n=1}^{\infty} a_n [\text{sinc}[\pi(na-k)] \cdot \\
 &\quad \cdot \frac{\bar{k} (i \sin(2\pi na K) + \cos(2\pi na K)) (\cos \Phi_{s_n} + i \sin \Phi_{s_n}) + d (-i \sin(2\pi na K) + \cos(2\pi na K)) (\cos \Phi_{s_n} - i \sin \Phi_{s_n})}{\bar{k}}] = \\
 &= \sum_{n=1}^{\infty} a_n \text{sinc}[\pi(na-k)] \left[\underbrace{\exp[i(2\pi na K + \Phi_{s_n})]}_I + \underbrace{\frac{d}{\bar{k}} \exp[-i(2\pi na K + \Phi_{s_n})]}_{II} \right] =: \sum_{n=1}^{\infty} A_{kn} a_n \tag{21}
 \end{aligned}$$

where we assumed b_0 to be zero and used $s_n =: a_n (i \sin \Phi_{s_n} + \cos \Phi_{s_n})$, a_n real.

APPENDIX B: SOME CONSIDERATIONS ON THE EVALUATION OF THE RECEIVER POSTERIOR PDF

Having rigorously marginalized out the stationary process g , we seek to maximize the posterior probability density function with respect to the template parameters. When projecting this equation into Fourier space, G is a diagonal matrix and the multiplication with \hat{f} transforms to a convolution that breaks this diagonality, leading to a coupling of the Fourier coefficients. While in the diagonal case the pdf completely factorizes in separate Fourier coefficients, which are independently and quickly integrated, there is only one N_{bin} -dimensional integral in the exact case. For typical N_{bins} of 1024 the computational demand is simply too high. Thus we decided for imposing diagonality in Fourier space, which discards the non-diagonal terms of f leading to a solvable integral. However, this discards phase information as now f has also the properties of a stationary process. Consequently we have derived an amplitude-only model. An Amplitude model alone does not contain valuable information about TOAs since the amplitude is invariant to time shifts. Therefore, we impose a model on the phase of the signal based on wrapped gaussians and being evaluated phenomenologically. This model uses the amplitude information to estimate the phase error of an observation at hand. The phenomenological way was chosen as the exact calculations for quadrature of the stochastic signal mix the phases of different Fourier channels in a highly non-trivial fashion again leading to an unwanted coupling of phases present only at the single pulse level.

This ansatz tries to find a compromise between numerical feasibility and accuracy. For the case of real measurements containing additional noise and fluctuations, we expected to grasp the uncertainties left with acceptable precision and lead to acceptable χ^2 -values at the end.

However losing the mathematical precision, we expect not to reach the accuracy of the classical timing codes when it comes to evaluating exact, nonfluctuating test data sets. Indeed our round mean square errors turned out to be 25% worse in these cases while the reached χ^2 was comparable accurate.

The question surely arises why we did not evaluate the pdf in time domain. In time domain, the main operator acting on the data is the inverse of a diagonal matrix (\hat{f}) plus a Toeplitz-matrix G ⁶ that can be further reduced to a circular Toeplitz-Matrix by the symmetry in \hat{f} . There exist fast ways (up to $\mathcal{O}(N \log N)$) of calculating the operators scalar product with the data (Ng & Pan 2010) however this does not reduce the dimensionality of the integration when it comes to varying the template parameters of interest. Unfortunately, expanding the operator linearly (to develop a mean field theory for the signal field, which should lead to a very good estimate and first order corrections for the correlations present in the problem) does not help in this case, as the derivatives refer to specific entries of the inverse matrix. Calculating the inverse matrix itself scales $\mathcal{O}(N^3)$ or best $\mathcal{O}(N^2)$ in our case, and as we need N entries for the calculation of the mean field gradient, the computational demand is again unacceptable. A class of transformations on f that leave the determinant of the operator invariant, combined with a pdf on the invariants of the transformation, could make exact evaluation of the derived equations possible. In lack of such a tool we decided to approximate the problem in Fourier space.

APPENDIX C: ITERATIVE SOLVER FOR FIELDS

Finding the maximum a-posteriori (MAP) probability set equals finding a global maximum for $\bar{\mathbf{f}}$ and $\sigma_{\mathbf{f}}$ both at the same time. Equations (16) and (17) explicate optimization w.r.t. one parameter only in the context of correct prior knowledge. Furthermore, we want to find the optimal solution w.r.t. both parameters. We may resolve these issues by assuming that there is only one maximum (which might not be the case as there might be several modes of radiation in the system etc.) and rely on a suitable iteration scheme. We implemented a method of steepest ascend. The Hamiltonian may be seen as a potential to be broken up in parts dependent on \bar{f}_k and σ_k for every Fourier coefficient. For the sake of shortness let us call this part $\Phi(\bar{\mathbf{f}}(t), \sigma(t))$. We now try to find the path of the steepest ascend, parametrized by t : $\gamma_t = (\bar{\mathbf{f}}, \sigma)$. Infinitesimally, isolines may be found using $d\Phi(\gamma_t) = 0$. Working this out for our Hamiltonian leads to

$$0 = \left\langle \frac{(\mathbf{f} - \bar{\mathbf{f}})^2}{\sigma^3} - \frac{1}{\sigma} \right\rangle \frac{\partial \sigma}{\partial t} + \left\langle \frac{\mathbf{f} - \bar{\mathbf{f}}}{\sigma^2} \right\rangle \frac{\partial \bar{\mathbf{f}}}{\partial t} \quad (\text{C1})$$

Consequently, the direction of the steepest ascend/descent follows the differential equation

$$0 = -\left\langle \frac{\mathbf{f} - \bar{\mathbf{f}}}{\sigma^2} \right\rangle \frac{\partial \sigma}{\partial t} + \left(\left\langle \frac{(\mathbf{f} - \bar{\mathbf{f}})^2}{\sigma^3} \right\rangle - \frac{1}{\sigma} \right) \frac{\partial \bar{\mathbf{f}}}{\partial t} \quad (\text{C2})$$

where orthogonality is easily shown. Since for most cases occurring setting $\bar{\mathbf{f}}$ according to (16) is an acceptable method while a reasonable update procedure of σ prior to optimizing it was missing, we update σ when updating $\bar{\mathbf{f}}$ by the method, but not vice versa. Integrating (C2) is approximated by

$$\Delta \sigma = \frac{\partial \sigma}{\partial \bar{\mathbf{f}}} \Delta \bar{\mathbf{f}} = \left(\frac{\langle (\mathbf{f} - \bar{\mathbf{f}})^2 \rangle}{\sigma} - \sigma \right) \frac{\Delta \bar{\mathbf{f}}}{\langle \mathbf{f} - \bar{\mathbf{f}} \rangle} \quad (\text{C3})$$

$$\Rightarrow \sigma_{new} = \frac{\langle (\mathbf{f} - \bar{\mathbf{f}}_{new})^2 \rangle + \langle (\mathbf{f} - \bar{\mathbf{f}}_{old})^2 \rangle}{2\sigma_{old}} \quad (\text{C4})$$

and determining $\bar{\mathbf{f}}_{new}$ as in (16). As is evident by derivation, this method is generic for a prior of the form of a Gaussian, however could be refined by also updating the $\bar{\mathbf{f}}$ value by similar considerations.

⁶ a matrix whose entries depend only on the distance to the diagonal